

Low-Cost, Portable, Multi-Wall Virtual Reality

Samuel A. Miller[†], Noah J. Misch[‡], and Aaron J. Dalton[§]

NASA Applied Sciences DEVELOP Program

Abstract

Virtual reality systems make compelling outreach displays, but some systems highly suitable for outreach, notably the CAVE, have design features that make using them for that purpose inconvenient. In the case of the CAVE, the equipment is difficult to disassemble, transport, and reassemble, and typically only large-budget research facilities can afford such a system.

We implemented a system like the CAVE that costs less than \$30,000, weighs about 500 pounds, and fits into a fifteen-passenger van. A team of six people have unpacked, assembled, and calibrated the system in less than two hours.

This cost reduction versus similar virtual reality systems stems from the unique approach we took to stereoscopic projection. We used an assembly of optical chopper wheels and commodity LCD projectors to create true active stereo at less than a fifth of the cost of comparable active stereo technologies.

The screen and frame design also harbor portability optimizations; the frame assembles in minutes with only two fasteners, and both pack into small bundles for easy and secure shipment.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and RealismVirtual reality I.3.2 [Computer Graphics]: Graphics SystemsDistributed/network graphics I.6.7 [Simulation and Modeling]: Simulation Support SystemsEnvironments

1. Introduction

1.1. Impetus

Communicating the significance of scientific data to a broad audience requires compelling visualization tools. DEVELOP, a student group within NASA's Earth Science Enterprise, demonstrates applications of NASA Earth observations and modeling at a number of stakeholder conferences each year. At such events, we traditionally presented animations of relevant data sets, but we came to find those lacking. The animations forced viewers to observe the data in one way, and videos playing on a projector did not communicate as compelling a message as we believed possible.

1.2. Our Approach

We created a low-budget, portable variant of the CAVE [CNSD*92] virtual reality system. We appreciated the immersion and interactivity the CAVE provides, but our budget did not cover a single major part of a traditional CAVE. Furthermore, a CAVE-like virtual reality device for use primarily at conferences suggested very different design constraints; we needed lightweight components that stored compactly and shipped reliably, and we wanted to assemble the system from its shipped state in a matter of hours. To meet those requirements, we designed afresh the stereo projection equipment and screen framework and implemented a conventional cluster of PC computer systems for rendering.

We used a cluster of commodity personal computers, one for each wall, as the imagery engines for our device. Advances in inexpensive graphics cards targeting computer games allow such a system to run applications that previously only specialized graphics workstations could han-

[†] e-mail:samuel.a.miller@nasa.gov

[‡] e-mail:noah@cs.caltech.edu

[§] e-mail:aaron.j.dalton@nasa.gov

dle. The system runs software frameworks such as NetJuggler [AGL*02], VRJuggler [Bie00], and CAVELib. We successfully run well-written software for frameworks such as CAVELib and VRJuggler without modifying them.

We replaced the traditional CAVE's expensive stereoscopic projectors with a novel assembly of commodity projectors, stepper motors, and optical choppers. We developed this system in two stages, at first to support stereo on a single display wall, and then to support it on multiple walls. It creates an active-stereo effect much like that of traditional systems, with some limitations that we detail. This document focuses on this stereo projection component.

Other components, such as the screen, frame, and mirrors, resembled those of traditional CAVEs conceptually but included incremental improvements, often upon their portability.

1.3. Other Virtual Reality Systems

Others have created different low-cost virtual-reality systems. The GeoWall [SDW02], popular in academia, has a single display screen and uses passive stereo. Belleman et al. developed a one-wall active-stereo virtual environment [BSdV01] using a PC and sync-doubler that split the output signal into left-eye and right-eye frames. Their paper also covered broadly the choices available to low-cost VR implementors.

Various commercial interests have also developed portable virtual-reality systems; however, for organizations such as ours, their prices are prohibitive. Fakespace Systems ROVR with its 15-minute setup time is very appealing. It utilizes active- or passive-stereo (user specified), but as with the systems above is limited to single-wall visualizations. Barco's passive-stereo Transportable CADWall is similar. Fakespace's Beacon projection system closely resembles ours in that it uses two projectors and an optical shuttering device to achieve active stereo imagery, and unlike those above, can function alone or in concert with other Beacon units.

There are various multi-wall virtual-reality systems specifically designed for mobility. Barco's Transportable I-Space system (with three to six walls) is active-stereo capable and is transportable, with a 1-2 day setup time. Iowa State's four-wall Baby Cave is reconfigurable from a small room (8-by-8-by-6 feet) into a 32-foot-long display and has a 2-hour setup time, but is limited to passive-stereo. The blue-C [GWN*03] immersive virtual-reality and 3D video-acquisition environment heavily utilizes liquid-crystal shuttering devices, but is designed for a stationary environment.

The system we present differs from those above in that it forms a multi-wall virtual environment for less than \$30,000, a price point well below traditional multi-wall VR systems.

We accomplished that primarily by diligently avoiding expensive specialized equipment.

2. Low-Cost Stereoscopic Projection

2.1. Overview

Very early in the design stage we chose to focus our efforts on active-stereo projection systems. This choice was motivated by the high cost of the polarization preserving rear-projection screen material, and the view-angle limitations imposed by typical polarization-based passive-stereo systems. Given these constraints, our greatest technical challenge lay in replacing the CAVE's stereo-capable projectors and frame-locked graphics cards. At first we tried to synchronize the frame refresh of our computers and drive a traditional stereo output device. Using the softgenlock [AGL*03] Linux kernel module and a TTL_PAPERS [DHM96] synchronization network, we achieved high-quality stereo across a group of four CRT computer monitors, each with a different computer driving it. Unfortunately, we could not extend this system to a visual immersion environment because that required replacing the monitors with projectors, and no projector with retail price below about \$20,000 refreshes its output quickly enough for immersive VR.

2.2. Single-Wall Implementation

In light of those projector constraints, we investigated a wholly different approach. An optical chopper wheel placed in front of a pair of commodity projectors with their images aligned, each displaying the image for a particular eye, creates an effect like that of conventional active-stereo projection systems. Instead of supplying a single graphics feed and rapidly switching between left- and right-eye images, we created two feeds, one per eye, and used the chopper to switch between exposing the projectors attached to each feed. Indeed, in terms of its implementation, the system resembles a passive stereo system; we drive the projectors in exactly the same way, and we use two aligned commodity projectors.

Although we learned of it after our system was implemented, our single-wall system is similar to Frohlich et al.'s [FHH*05] multi-user stereo environment, particularly with regard to mechanical projector-shuttering and stereo-eyeglass synchronization. Both our system and that described in [FHH*05] have a historic foundation in Hammond's 1924 patent [Ham24], and a more recent European patent [Pal02] proposing a related idea. Fakespace's Beacon system mentioned earlier employs a similar method, but uses a pair of liquid-crystal shutters instead of the mechanical shuttering employed in our design. Although this methodology has many merits, the high ambient-light conditions present at our target venues, combined with the inefficient light transmission and high cost of commercially available

liquid crystal shutters, made the more cumbersome chopper/motor arrangement preferable in all regards.

We aligned the projector's images using a locally-developed wood alignment box with numerous adjustments; the chopper wheel was an aluminum disc twelve inches in diameter and 1/8" thick with two ninety-degree openings. A small DC motor attached to the face of the alignment box drove that optical chopper. We selected inexpensive liquid-crystal display projectors with XGA resolution; SXGA-resolution projectors would improve our overall visual quality but raise our projector costs roughly threefold.



Figure 1: *Early alignment box with aluminum chopper wheel.*

Users viewed the system through the same liquid-crystal shutter glasses used in standard CAVEs. One normally controls such glasses with infrared emitters that attach to the computer's graphics card. Since the optical chopper, not the graphics card, dictated switches between the left- and right-eye frames in this system, we instead placed a reflective sensor [Fai00] next to the chopper, along the top of one projector lens. The sensor allowed current to flow only when a reflective surface, such as a closed portion of the chopper, passed, so the circuit created the square-wave signal the emitter required.

2.3. Multi-Wall Implementation

That assembly of two projectors and an optical chopper created a high-quality active-stereo display on a single wall, but

extending the technique to four walls required introducing infrastructure to synchronize the optical choppers. Otherwise, each wall would switch the exposed eye in an uncoordinated fashion and users would not see the entire display in stereo.

We examined a range of motors that could operate in a phase-synchronized fashion while turning at the speed of the DC motor in the one-wall system. Servo motors and stepper motors suited the application best, the former utilizing closed-loop control, and the later open-loop. We selected a stepper-motor system because it met our requirements at one-fifth the cost of comparable servo-motor implementations. Because stepper motors rotate a given number of degrees for each pulse sent to their drivers, multi-wall synchronization is achieved by aligning the choppers on each projector box initially, then triggering the steppers in parallel with a simple clock signal.

To allow us to use a smaller stepper motor, we switched from the cut-aluminum chopper wheel to one we cut from a sheet of 7-mm hydroponics Mylar. This reduced costs by eliminating a dependency on machined parts, and it improved system safety. Whereas the aluminum disk would severely injure a person should it contact a body part, touching the moving Mylar wheel simply tears the wheel.

The complete system uses the projectors and glasses of the one-wall prototype and synthesizes the emitter drive signal in the same way. Since a step-clock controls the stepper motors, one could instead derive the emitter signal from that. We chose not to do this since it involved fabricating electronics with which we had limited familiarity, and because it offered no great advantage.

2.4. Stereo Functional Comparison

Despite its structural similarity to passive stereo systems, output from our system is pure active stereo and practically indistinguishable from that of other active-stereo systems. One can distinguish between the two based upon a parameter we termed "single-image exposure time," the fraction of time during which the actual light the configuration emits represents the image for exactly one eye. The greater the single-image exposure time, the higher the theoretical output quality.

An active-stereo system of a computer and a CRT monitor has almost complete single-image exposure; the image flips between the left- and right-eye images nearly instantly. Depending upon the size of the optical chopper wheel and the number of cut-outs in it, single-image exposure time in this system can fall as low as 70%. In practice, we could not observe a subjective difference in stereo quality across the range of single-image exposure times we tested.

With this system, one can select virtually any refresh rate

by adjusting motor speed; in practice a 72Hz refresh rate result in a flicker-free image. The stereo refresh rate is simply equal to the product of the motor speed and the number of pairs of open and closed segments in the chopper wheel. We used a wheel with two such pairs at first but later switched to a four-pairs wheel. Although fewer pairs implies higher single-image exposure time, we found that image quality was not highly sensitive to that parameter. Consequently, we were able to use more pairs and therefore turn stepper motors at lower speeds for a given refresh rate. Appropriately selected motors and chopper wheels allow a wide range of refresh rates: a distinct advantage over traditional active-stereo implementations.

2.5. Limitations

The motors in this stereo system generate a moderately high-pitched noise while operating, and the alignment frame to which they attach amplifies that noise. We reduced the latter effect considerably during development by using plastic screws and plastic insulation at connections between the motor bodies and the framework, but the system remains considerably noisier than the single stereo-capable projectors traditional CAVEs use. At the typical 72Hz operating point the Mylar chopper wheels' noise output is similar to a desktop computer, however they become louder and tend to vibrate when the stereo refresh-rate is increased beyond about 75-80Hz. A more rigid wheel material would mitigate the vibration, and appropriate enclosures would significantly reduce the noise output.

The standard shutter-glasses our users wear use polarizing filters to occlude the eyes, and that interacts destructively with the LCD projectors, which emit polarized light. A user will see distorted color in his or her peripheral vision always and everywhere if he or she tilts his or her head to one side. Tilting in one direction, the pictures look red, and tilted the other direction, the image looks blue. In some respects, this resembles the sort of distortion one experiences in a passive stereo system, in which tilting one's head eliminates eye separation and exposes both images. In this case, the stereo effect remains intact, but the glasses filter out some components of the light, reducing it to one primary color.

We investigated replacing the LCD projects with single-chip digital light-processing projectors, which offer higher contrast ratios at slightly higher price points than LCD projectors of similar pixel resolution and brightness. Notably, these DLP projectors emit un-polarized light, so shutter glasses do not distort their output as they do that of LCDs. We discovered a different problem with them during experiments, though; DLPs interact destructively with the optical chopper. Internally, the DLP projector has a color wheel [Tex03], and the projector exposes an image in each color in turn, very rapidly, so humans perceive true color. Placing an optical chopper wheel in front of this output defeats that technique;

the chopper blocks some wedge of the projector's color output, washing out a particular primary color in part of the image. The washed-out color cycles with the beat frequency between the projector and the chopper. Theoretically, one could eliminate this problem by selecting a chopper angular speed such that the number of times it exposes one eye's image each second is an integral multiple of the number of times the projector color wheel cycles through all colors each second; we did not pursue that possibility further.

As three-chip DLP projectors fall in price, their costs may rival that of the projectors in this assembly. At present, though, they typically cost in excess of \$20,000, and therefore a stereoscopic projection system based upon them costs at least five times what this system does. The economics are more in favor of three-chip DLPs for SXGA resolution; in that case for our application the DLP projectors are only two to three times the cost of typical LCD SXGA units.

3. Screen, Frame, and Mirrors

The frame upon which we hang the projection screen is constructed of two-inch square aluminum tubing; welded aluminum joints slide snugly into the tubing to form the corners. The frame forms a box 10 feet on each side and 8-feet tall, mirroring the projector aspect ratio. Using this combination of straight tubes and simple corner pieces, the frame is extremely rigid, yet uses no fasteners. No-fastener construction allows the frame to be set up and taken down rapidly without tools, and the aluminum construction makes the frame light-weight.

For a number of reasons, we did not accommodate suspending a projector and mirror overhead for projecting onto the floor, as CAVEs usually do. The ceiling clearance in our engineering space did not allow for such an assembly, and many of our outreach venues also lack the ceiling height overhead projection requires. Furthermore, designing a frame that suspends a mirror over a human would have extended our initial development cycle substantially through increased safety-related verification requirements. For purposes of this initial prototype we opted to focus instead on the stereo-related issues.

Minimizing the visible seams between the rear wall and the two side walls contributes greatly to the sense of immersion users experience. We used taut wires (Figure 2) to define those interfaces, but we avoided bowing in the wire by stretching the screen minimally along the horizontal axis. The wires that form the corners are tensioned using threaded inserts at one end (Figure 3). The nuts on the inserts are the only fasteners in the entire frame!

A single piece of screen material forms all three walls of the virtual environment; it clings to the frame using zippered pockets sewn onto each edge. Figure 3 shows the black fabric of the pocket and zipper. This design facilitates rapid setup and uniform screen tensioning.




Figure 2: Interior edge. Mars Pathfinder image courtesy NASA/JPL-Caltech.




Figure 3: Corner detail.

As in the CAVE, mirrors fold the throw distance of the projectors, thereby constraining our overall system footprint to approximately 24'x24'. We use rear-surface plate-glass mirrors designed for decorative use in homes and offices in place of more traditional Mylar mirrors. Though cumbersome, they cost less and withstand shipping better than do Mylar mirrors.

4. User Input and Software Applications

Because our low-cost virtual environment is intended to function as an outreach tool for NASA's earth-science geospatial data demonstrations, users need to navigate through expansive environments. This type of navigation is particularly well-suited to standard computer-gaming input devices, which - although they cannot track the user's gestures and movement - do allow the user to easily change direction and velocity within the virtual world.

Additionally, a typical interaction scenario for our devices involves multiple users viewing the virtual environment simultaneously. In such situations head-tracking only benefits one viewer. Thus, for our target applications and audiences, tracked user interfaces for head and hands are superfluous: we chose instead to use a wireless gamepad as our primary input device.

5. Summary

The combination of this approach to stereoscopic projection, use of commodity personal computers for rendering, and ancillary optimizations, enabled us to construct a portable version of the CAVE for under \$30,000. The system weighs less than 500 lbs and fits into the back of a fifteen-passenger van (Figure 4) with one remaining row of seats. See Table 1 for weights and costs of each major system component.



Figure 4: Mirrors (boxed), frame members, and alignment boxes in van.

We displayed the one-wall prototype at the 2003 Southern Growth Policies Board Annual Conference. Later, we assembled the three-wall system in a conventional conference room at NASA Headquarters. In the latter case assembly took six people less than two hours, including time for unpacking, assembling, and alignment calibration.

Table 1: Weight and Cost of Major System Components

Component	Quantity	Weight (lb, each)	Cost (\$, each)
Computer	5	30	1500
Projector	6	7	1800
Alignment Box	3	30	250
Mirror	3	30	100
Frame	1	80	900
Screen	1	40	3000
Emitters	6	1	200
Glasses	4	1	425
Other	1	20	300
TOTAL		522	26450

6. Acknowledgements

Our thanks to Ronald Birk and Martin Frederick of NASA Headquarter's Science Mission Directorate for funding this work, and to NASA's DEVELOP program for empowering students to pursue interesting projects. Additionally, special thanks to Michael Ruiz for facilitating DEVELOP's activities, and to Dr. John Quagliano and Thomas Spencer for their advice on the project.

References

- [AGL*02] ALLARD J., GOURANTON V., LECOINTRE L., MELIN E., RAFFIN B.: Net juggler : Running vr juggler with multiple displays on a commodity component cluster. In *IEEE Virtual Reality Conference (2002)*, Virtual Reality Proceedings, Annual Conference Series, IEEE, IEEE, pp. 273–274. 1
- [AGL*03] ALLARD J., GOURANTON V., LAMARQUE G., MELIN E., RAFFIN B.: Softgenlock: Active stereo and genlock for pc cluster. In *Proceedings of the Joint IPT/EGVE'03 Workshop* (Zurich, Switzerland, May 2003). 2
- [Bie00] BIERBAUM A.: *VR Juggler: A Virtual Platform for Virtual Reality Application Development*. Master's thesis, Iowa State University, 2000. 2
- [BSdV01] BELLEMAN R. G., STOLK B., DE VRIES R.: Immersive virtual reality on commodity hardware. In *ASCI 2001 Conference* (Heijen, the Netherlands, May 2001), Advanced School for Computing and Imaging. 2
- [CNSD*92] CRUZ-NEIRA C., SANDIN D., DEFANTI T., KENYON R., HART J.: The cave: Audio visual experience automatic virtual environment. In *Communications of the ACM* (1992), vol. 35, ACM, ACM Press, pp. 65–72. 1
- [DHM96] DIETZ H. G., HOARE R., MATTOX T.: A fine-grain parallel architecture based on barrier synchronization. In *Proceedings of the International Conference on Parallel Processing* (August 1996), pp. 247–250. 2
- [Fai00] FAIRCHILD SEMICONDUCTOR: *QRD113/1114 Reflective Object Sensor*, 2000. 3
- [FHH*05] FROHLICH B., HOCHSTRATE J., HOFFMANN J., KLUGER K., BLACH R., BUES M., STEFANI O.: Implementing multi-viewer stereo displays. In *Proceedings of the WSCG 2005 Conference* (Plzen, Czech Republic, 2005), WSCG'2005, UNION Agency - Science Press. 2
- [GWN*03] GROSS M., WURMLIN S., NAEF M., LAMBORAY E., SPAGNO C., KUNZ A., KOLLER-MEIER E., SVOBODA T., GOOL L. V., LANG S., STREHLKE K., MOERE A. V., STAADT O.: blue-c: A spatially immersive display and 3d video portal for telepresence. In *ACM Transactions on Graphics* (2003), vol. 22, ACM, ACM Press, pp. 819–827. 2
- [Ham24] HAMMOND L.: Stereo-scopic motion picture device (us patent 1506524), Aug. 1924. 2
- [Ols02] OLSON E.: *Cluster Juggler - PC Cluster Virtual Reality*. Master's thesis, Iowa State University, 2002.
- [Pal02] PALOVUORI K.: Apparatus based on shutter function for projection of a stereo or multichannel image (eu patent wo03003750), June 2002. <http://v3.espacenet.com/textdoc?DB=EPODOC&IDX=WO03003750>. 2
- [SDW02] STEINWAND D., DAVIS B., WEEKS N.: Geowall: Investigations into low-cost stereo display systems. USGS Open Field Report, 2002. <http://geowall.geo.lsa.umich.edu/papers/Geowall.pdf>. 2
- [SFLS00] SAMANTA R., FUNKHOUSER T., LI K., SINGH J. P.: Hybrid sort-first and sort-last parallel rendering with a cluster of pcs. In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS Workshop on Graphics hardware* (Interlaken, Switzerland, August 2000), ACM, ACM Press, pp. 97–108.
- [Tex03] TEXAS INSTRUMENTS: *How DLP Technology Works*, 2003. http://www.dlp.com/dlp_technology/dlp_technology_overview.asp. 4